

Seamlessly Integrate With OpenStack, Increasing Efficiency and Reducing Operational Costs

OpenStack, commonly referred to as "the Linux of the Cloud," allows companies to utilize open-source initiatives and a transparent and collaborative approach of implementing public and private cloud solutions, enabling business agility, infrastructure elasticity and operational simplicity. While OpenStack provides a flexible framework, it's important that the cloud infrastructure, composed of compute, network, and storage resources, runs at maximum performance and efficiency to assure application performance. Dell EMC server and storage, combined with Mellanox® networking, provide the ideal hardware architecture to properly support web-scale IT with Red Hat OpenStack as the cloud management platform.

EFFICIENT CLOUD DEPLOYMENT

Leveraging disaggregation and virtualization technologies required for efficient cloud computing comes at a significant cost manifested by performance penalties such as slow data communication and storage access, and heightened CPU utilization. To address these challenges, Dell EMC partners with Red Hat and Mellanox to bring to market a highly efficient, hardware-accelerated, and tightly integrated cloud data center solution. Combining open standards-based Dell EMC compute, networking, storage and management tools with Red Hat OpenStack Platform (RHOSP) 13 allows for rapid delivery of IT services for not only Communication Service Providers (CSPs/Telcos), but also Hosting Service Providers, and Managed Service Providers (MSPs) as well as high-end enterprise private cloud builders.

The solution allows service providers to remove proprietary infrastructure, decrease total cost of ownership and automate operations, as they deliver new sets of services and applications to end customers. Proprietary equipment is replaced with virtual machines (VMs) running on Dell EMC 14th generation PowerEdge Servers, RHOSP 13 with Mellanox ConnectX® intelligent networking adapters running offload technologies. This reduces OpEx and enable a quicker time to revenue.

SERVICE ORIENTED, KINETIC INFRASTRUCTURE

Kinetic infrastructure treats networking, storage, and compute as fluid resource pools that can be composed on-the-fly. This includes full deployment automation to configure the hardware and software and, simplified management. Capacity can be changed based on software applications requirements and makes a data center more scalable and flexible. This means companies can reduce stranded assets, over provisioning and improve efficiency.

SR-IOV I/O VIRTUALIZATION

Single Root I/O Virtualization (SR-IOV) allows a network adapter, to provision access to its resources among various PCIe hardware functions. This allows traffic streams to be delivered directly between the virtual machines and their associated PCIe partitions, giving applications direct access to the I/O hardware. As a result, the I/O overhead in the software emulation layer is eliminated. SR-IOV enables VMs to achieve network performance that is nearly the same as in non-virtualized environments.

Mellanox NICs support basic SR-IOV as well as advanced features such as SR-IOV High Availability (HA) and Quality of Service (QoS). SR-IOV HA provides a redundancy mechanism for Virtual Functions (VF) by using Link Aggregation

KEY BUSINESS BENEFITS:

- Out-of-Box End-to-End Cloud Management Platform
- Pre-validated Reference Architecture for NFV Deployments
- Optimization of OpenStack Cloud through Software Accelerators and Hardware Offloads Including:
 - SR-IOV PCIe I/O Virtualization
 - OVS DPDK
 - Overlay Networks
 - Stateless Offloads
- Highest DPDK Bare Metal Packet, Rate at Over 148Mpps to Build VNFs
- 55Mpps Through ASAP² with Zero CPU Utilization
- Lower Capital and Operating Costs
- Red Hat Open Source Software Combined with Dell EMC Compute and Storage and Mellanox High-Performance and Efficient Networking Technologies

Modular infrastructure that is bringing compute, storage, and networking together!

Group (LAG) to bind two VFs from two different ports on the same NIC together, and exposes the bundle as one VF to the VM. When one VF in the bundle fails, the other VF continues forwarding traffic without affecting VM I/O operations.

THE BEST PERFORMANCE FOR DPDK

Data Plane Development Kit (DPDK) reduces overhead caused by interrupts that are sent each time a new packet arrives for processing. DPDK implements a polling process for new packets to achieve the key benefits of significantly improving processing performance while eliminating PCI overhead and maintaining hardware independence. Although DPDK technology consumes CPU cycles, Mellanox ConnectX-5 adapters offer the industry's highest bare-metal packet rate of 148 million packet per second for running OVS or VNF cloud applications over DPDK.

ASAP²

Accelerated Switching and Packet Processing (ASAP²) was developed by Mellanox as an open and high-performance vSwitch and vRouter offload technology. ASAP² offers hardware acceleration through SR-IOV data path (fast-path) for high-throughput flows along with unmodified standard OVS control path for SDN flexibility and programming of match-action rules. ASAP² offloads networking functions such as overlay tunneling, routing, security and load balancing to the Intelligent NIC's embedded switch (e-switch).

Performance tests conducted in Red Hat's lab verified ASAP² delivered close to 100G line-rate throughput for large VXLAN packet encap/decap without consuming any CPU cycles. ASAP² also boosted vanilla OVS performance by 20x and OVS-DPDK performance by 8x to 55 million packets per second with zero CPU utilization for 64-byte VXLAN packets. ASAP² frees CPU cores for deploying more cloud-native applications and VNFs on the same server and is fully integrated with RHEL 7.5 and RHOSP 13.

RDMA/ROCE TO OVERCOME INEFFICIENCIES

The large overhead associated with stateful protocols such as TCP dictates that it is not an ideal transport

protocol for software defined scale-out storage applications, especially when storage media gets faster transitioning from hard disks to solid-state drives (SSD) and to Non-Volatile Memory Express (NVMe).

Remote Direct Memory Access (RDMA), on the other hand, is a protocol designed for high-speed links within data center environment that can overcome the inefficiencies of TCP. RDMA usually runs over Converged Ethernet (RoCE). RDMA implements a kernel bypass to read and write network semantics and full transaction offload to RDMA capable NIC devices to guarantee the highest possible throughput, lowest latency, and minimal CPU overhead, making it ideal for storage access. Typically RDMA over Converged Ethernet (RoCE) requires the network to be configured for lossless operation, however, Mellanox has recently enhanced RoCE with built-in error recovery mechanisms. While a lossless network has never been a strict requirement, customers typically configure their networks to prevent packet loss and ensure the best performance. With this new version, RoCE can be deployed on an ordinary Mellanox Ethernet network. By utilizing RDMA or RoCE, virtual servers can achieve much higher I/O performance because the majority of packet processing is offloaded to the NIC. This further enables increased performance, improved latencies and significantly reduced CPU overhead. The net effect is an improvement of overall server and application efficiencies.

CONCLUSION

The Dell EMC Ready Architecture for Red Hat OpenStack Platform is a pre-validated, proven, adaptable reference architecture for NFV deployments. As IT organizations transition to cloud-based and service-centric infrastructures, the need to gain network and server efficiencies is paramount to transition beyond 10Gb server I/O. With Dell EMC cloud infrastructure hardware, Red Hat OpenStack Platform and key technologies from the Mellanox adapter, users are able to accelerate virtual and bare-metal networks and reduce CPU utilization through hardware-based offloads for increased scalability, greater flexibility and the highest efficiency in the modern software-defined data center.

WANT TO LEARN MORE?

Dell EMC Ready Architecture for Red Hat OpenStack Platform Architecture Guide:

https://www.dell.com/resources/en-us/asset/technical-guides-support-information/solutions/dell_emc_red_hat_ready_architecture_guide_v13.2.pdf

Dell EMC Ready Architecture for Red Hat OpenStack Platform Solution Brief:

https://www.dell.com/resources/en-us/asset/brochures/solutions/h17453_dell_emc_ready_architecture_for_red_hat_openstack_platform_solution_brief.pdf



**Accelerate Network Performance and Lower Costs
with 25/100GbE Networking from Mellanox**